

# Research Proposal

Daniel Genin

July 21, 2006

## Introduction

Our world is full of large and complex networks both natural and man-made : communications networks, power grids, trophic webs, social networks are but a few examples. Yet our understanding of these important networks is primitive. The situation is somewhat similar to the state of metallurgy in the 16th century : while sophisticated steel weapons were already being made by purely empirical methods, and gave great advantages to the nations producing them, the understanding of the metallurgical processes was utterly lacking [10]. The development of atomic-level metallurgy in the 20th century enabled creation of high-strength and light-weight materials which in turn revolutionized several manufacturing industries (e.g. aerospace). True understanding is thus key to progress. While in most sciences great progress has been made in analyzing individual systems there are virtually no general approaches for analyzing the intricate behavior that arises when large numbers of such systems are linked into complex networks.

The aims of network analysis vary with context but one broad goal is the discovery of mathematical laws describing spatial-temporal evolution of quantities of importance associated with nodes and links [11]. These could be, for example, capital for business networks, bandwidth utilization for communications networks, and concentrations of chemicals for metabolic networks. Development of a general theoretical framework for such analysis remains a task for the distant future; nevertheless, incremental progress toward is being made at least in some disciplines.

One such discipline is the study of communications networks, including the Internet, which has been suggested as the perfect test ground for network science [4]. While it is manifestly a large and complex, we know everything about the design and functioning of its components, further, any theoretical prediction can be tested with copious measurements [18]. A good illustration was discovery of the origin of self-similarity in IP-level traffic. Initially self-similarity was attributed to self-organized criticality of the Internet [9],[13]. This was a natural conjecture because many critically self-organized systems exhibit scale-free distributions. This, however, was not born out by measurements of the actual Internet traffic [16]. Ultimately self-similarity of traffic was found to result from

heavy tailed distribution of session sizes at application level combined with a Poisson distribution of session arrival times. The heavy tailed-ness of session sizes has now been repeatedly confirmed by Internet measurements [3],[17].

The main contribution of our project will be to make a step in the direction of the grand goal – deriving laws for spatio-temporal evolution of networks – in the context of the Internet. More specifically, building on the current success of fluid approximation models, we will study behavior of large Internet-like communications networks with a variety of topologies under a variety of traffic mixes and loads.

## Application and Relevance

It is a testament to the vision of engineers of the original DARPA Internet that their design was able to withstand with little changes the dramatic modifications that the Internet underwent in the intervening 30 years. It survived exponential growth in the number of connected hosts [18], incorporation of new physical carriers, and arrival of numerous new and demanding applications.

While we have been successful thus far, the design of sophisticated networks like the Internet can not indefinitely rely on the intuitive insight of engineers and scientists. As complexity of crucial infrastructure networks increases, it becomes harder and harder to anticipate problems that may arise as the result of unforeseen collective behavior. The warning signs of this emerging trend have already started to surface. Among the more recent and noticeable was the cascading power outage that affected much of the Northeast US in August of 2003. The Internet also suffered a "congestion collapse" in October of 1986 [4], following which the TCP protocol was augmented with a congestion avoidance algorithm to become the protocol that we know today. There is danger for complex networks in general and for the Internet in particular, that without a theoretical framework for analyzing collective behavior a patchwork of solutions will emerge each fixing some shortcomings but always creating new vulnerabilities. It is thus crucial that tools for studying behavior of complex networks be developed in the nearest future.

## Background

Recently progress has been made in modeling aggregate behavior of congestion control mechanisms. The main idea is to consider the task of congestion control as a distributed optimization problem [12]. Consider a collection of routes indexed by a set  $R$  and links indexed by  $L$ . Each link  $l$  has a finite capacity  $c_l$ . Each route is composed of a number of links. We will write  $l \in r$  for  $l \in L$  and  $r \in R$  if the link  $l$  belongs to the route  $r$ . This information can be summarized in a matrix  $P$  which has a 1 in the  $l$ -th row and  $r$ -th column if  $l \in r$  and 0 otherwise. To each route  $r$  associate a source with transmission rate  $x_r(t)$ . The aggregate load on the link  $l$  is given by  $y_l = \sum_{r:l \in r} P_{lr} x_r$ , or

$\mathbf{y} = P\mathbf{x}$  in vector form. Each link  $l$  has a congestion measure  $-p_l(t)$ , which can be taken as the packet marking probability. Cumulative congestion signal for route  $r$  is  $q_r = \sum_{l:l \in r} P_{lr}p_l$ , or  $\mathbf{q} = P^T \mathbf{p}$  in vector notation. Source rates are taken to be governed by  $\frac{d}{dt}x_r(t) = f(x_r(t), q_r(t))$  and congestion measures by  $\frac{d}{dt}p_l(t) = g(p_l(t), y_l(t))$ . Arbitrariness in the choice of  $f$  and  $g$  gives this approach flexibility. By choosing different functions it is possible to approximate existing end-to-end congestion avoidance algorithms (e.g. TCP) and queue management algorithms (e.g. RED or ECN). This formalism allows to study the aggregate dynamics of congestion control algorithms prior to their deployment, which should allow to avoid unpleasant surprises, like the 1986 "congestion collapse".

To complete this setup assume every source has a utility function  $U_r(x_r)$ . The goal of congestion control can then be expressed as the following optimization problem :

construct  $f$  and  $g$  so that

$$U = \sum_{r \in R} U_r$$

is maximized subject to

$$y_l \leq c_l$$

for all  $l$ .

Considering  $x_r$  as primal variables and  $p_l$  as dual variables we have the following Lagrangian for this optimization problem

$$L(x, p) = \sum_r U_r - \sum_l p_l(y_l - c_l) \quad (1)$$

$$= \sum_r (U_r - x_r \sum_l P_{lr}p_l) + \sum_l p_l c_l \quad (2)$$

The dual problem can be written as

$$\min_{p_l \geq 0} \max_{x_i} \sum_r (U_r - x_r \sum_l P_{lr}p_l) + \sum_l p_l c_l \quad (3)$$

Thus  $U$  is maximized when individual net utilities  $U_r - q_r x_r$  are maximized, provided that prices,  $p_l$ , are set to the Lagrange multipliers. Observe that net utilities can be maximized by individual sources since they depend only on the cumulative congestion signal  $q_r$ . It is also possible to take  $\dot{p}_l$  depending only on  $y_l$  so that in equilibrium  $p_l$  will be equal to the Lagrange multipliers. Hence the problem of congestion control can be solved by sources maximizing their individual net utilities provided that  $f$  and  $g$  are chosen appropriately.

Here one can glimpse promise of a general theory of complex networks. The model just described parallels ecological models of populations competing for limited resources. A dictionary can be constructed in which sources and source rates above correspond to populations and population growth rates, the links and their capacities – to limited resources and their environmental capacities, the congestion feedback signal – to death rate. Thus results from ecology of

competing populations can be translated into results about networks of sources responding to congestion signals and vice versa. Even more, one may hope that the knowledge accumulated in these two disparate areas can be aggregated and distilled into a more general theory properly belonging to the field of complex networks.

The above framework, sometimes referred to as fluid approximation, has been extensively used in a sequence of papers beginning with Kelly et al. [7] to construct progressively more realistic models of Internet traffic with TCP-like congestion control. In [7] it was proved by constructing Lyapunov functions that if delays in propagation of signals are ignored then under some mild regularity conditions on  $g$ , the TCP-like congestion control scheme has a globally asymptotically stable equilibrium. This work was first extended to the case with delays but identical round trip times by Johari and Tan [5], and then to the case with heterogeneous round trip times by Massoulié [8] and Vinnicombe [15]. Introduction of propagation delays complicated the analysis because dynamics is governed by a system of delay differential rather than ordinary differential equations, and Lyapunov functions can not be easily constructed. Instead the authors of the later publications used methods of linear stability analysis and control theory to prove local asymptotic stability of the equilibrium, provided that  $\langle \text{gain} \rangle \times \langle \text{round trip time} \rangle \leq B$  and  $y \frac{d}{dy} \log(g) \leq B$  at equilibrium for some constant  $B$ , for all sources. Interestingly for the usual TCP algorithm stability is guaranteed only if the number of packets "in flight" is large enough and a modification of TCP free from this defect had been proposed [15].

## Proposed Research

The proposed research follows three main directions :

1. Extend the fluid approximation model for TCP traffic to application level protocols, e.g. HTTP;
2. Explore influence of network topology and changing demand from sources on dynamics of the fluid approximation model;
3. Validate results by forming testable hypotheses and verify them against simulations and existing data on Internet traffic.

We describe each of the directions in turn.

### Application level fluid approximation protocol model

The fluid approximation model approximates TCP dynamics over relatively short time-scales comparable with route round-trip times [6], the major cause being that the sources' demands are assumed to be constant. On longer time scales this assumption will be violated as demands of a source will vary even in the course of a single TCP connection, for example, as the user moves from

reading news, to browsing images, to streaming video from the same site. Yet it is a perfect platform on which to build longer time-scale models.

A key feature of complex systems is hierarchical structure, where modeling higher hierarchical levels often requires only a rough approximation of the lower levels. A good example is the two-leveled hierarchy of quantum mechanics and chemistry. They are largely insulated from each other by the difference in binding energies. So some quantum mechanical aspects of a molecule need be modeled only approximately to model its chemical properties. For instance, structure of the atomic nucleus, aside from charge, is irrelevant to chemistry. Similarly approximation of short time-scale TCP dynamics can become a plug-in module in models of higher level protocols operating on longer time-scales.

Success of the distributed optimization approach warrants an attempt to extend this line of research to application level protocols sitting above TCP, as already suggested in [4]. To date research on application level protocol traffic has focused mainly on simulation rather than analysis. I propose the following analytical model for the study of application level protocol traffic, extending the fluid approximation model.

Taking fluid approximation model as a basis, extend it by equipping each source with an exponentially distributed stochastic alarm clock, i.e. a Poisson point process with rate  $\lambda_r$ . Alarm signals will correspond to initiation of application sessions, e.g. web browsing. When the alarm of a source rings at time  $t$ , the source starts a TCP session by selecting a session size  $l_t$ , according to distribution  $L$ , and a utility function  $U_l$  from a one parameter family, determined by session size. The source can also select a route randomly from some distribution. A session continues until time  $u$  such that  $\int_t^u x(t)dt = l_t$  at which point the source ceases and resets its alarm. In the intervals between rings the source rates evolve according to an approximation of the dynamics obtained from the model of TCP. For example, one could assume that the rates evolve linearly toward the equilibrium corresponding to the particular combination of the utility functions or the full complexity of the underlying TCP dynamics can be incorporated by allowing the rates to evolve according to the model's system of differential equations.

The proposed model retains the flexibility of the distributed optimization approach to modeling TCP. Model parameters such as session size distribution, alarm clock rate distribution and utility function family, can be adjusted to model various application level protocols. Traffic types can be modeled by separating sources into classes, each endowed with its own parameters.

Dynamics of this model resembles a random walk driven by alarm signals. In recent years substantial body of work has been accumulated on systems of this type under the general title of random dynamical systems [1]. Using these methods I should be able to extend at least the stability results from TCP studies.

## Network dynamics as a function of changing demand and network topology

The role of network topology in the distributed optimization approach has so far received no attention. Existing results were derived for networks with arbitrary topology. While this indicates robustness of the TCP-like congestion control, the convergence rates, the path to equilibrium and possibly even uniqueness of the equilibrium do depend on network topology. Network topology will play an even more important role in traffic dynamics under shifting load. For example, a sudden rise in demand, like a flash crowd, in one part of the network may propagate through the network causing decrease in source rates across the network, in turn destabilizing TCP congestion avoidance.

To study the role of network topology in the fluid approximation model we must first identify parameters of topology affecting the model most directly. One example is the *link route-degree*, which can be defined as the function  $d : L \rightarrow \mathbb{N}$ , where

$$d(l) = |\{\text{the number of routes using } l\}|.$$

In terms of the routing matrix  $P$  used in the model,  $d(l)$  is just the sum of the entries in the  $l$ -th row. From  $d$  we can obtain the *link route-degree distribution*,  $D$ , which counts the number of links with a particular value of  $d$ . In the context of fluid approximation model this is a particularly meaningful measure of network complexity because it corresponds directly to the degree of coupling in the system of differential equations. For example, if  $D(k) = 0$  for all  $k > 1$  then all the routes are disjoint and the system of equations decouples completely. The greater dispersion in distribution  $D$ , the greater is the coupling between the differential equations.

To study the affect of this and other such quantities we will begin with a time-independent model, i.e., one in which sources' demands do not change with time, and study its behavior first in linear approximation near equilibrium and then, using a combination of analytical and numerical methods, on the whole phase space, for a variety of realistic network topologies like the recently developed "small world" and scale-free network models [14]. The results obtained here can also be used to assess the affects of topology on the proposed application level protocol model. The next, more interesting step, will be to study how the traffic load migrates across the network with varying source demands. Here various scenarios such as a gradual increase in demand, a flash crowd, or a DDOS attack will be analyzed analytically, to the extent it is possible, and simulated for a variety network topologies. As exact solutions will be impossible to obtain analysis will concentrate on identifying symmetries that may restrict dynamics and gross dynamics characterizations (e.g. oscillating, monotonic or chaotic). The results of this study should shed light on the important question of how the Internet might behave under stress as well as what topologies are most robust and resilient to various disturbances.

## Validation of results

Finally, to close the loop so to speak, the results from the proposed work will be tested using simulations and existing collections of Internet traffic data.

Computer simulations are particularly useful because the model, the environment and the inputs can be completely controlled. There are a number of platforms for simulating large-scale networks [2]. A particular approach developed by Yuan and Mills based on cellular automata [19], has the advantage of being able to simulate a large number of nodes over extensive time periods. Other approaches use fluid approximation and discrete event simulation but have the disadvantage of being restricted to simulating relatively few hosts. Depending on the hypotheses to be tested one or the other of these approaches may be appropriate. In addition, there is voluminous publicly accessible data on Internet traffic e.g. CAIDA. Although, one does not have control over measurement conditions, this data may still be used to test hypotheses predicting particular traffic characteristics. Thus we can not only form conjectures on the basis of our model exploration but also verify the validity of these predictions.

## Qualifications

My background in the theory of dynamical systems, ergodic theory and computing puts me in a very good position to successfully complete the proposed work.

## References

- [1] Ludwig Arnold. *Random Dynamical Systems*. Springer, 2003.
- [2] Lee Breslau, Deborah Estrin, Kevin Fall, Sally Floyd, John Heidemann, Ahmed Helmy, Polly Huang, Steven McCanne, Kannan Varadhan, Ya Xu, and Haobo Yu. Advances in network simulation. *IEEE Computer*, 33:59–67, 2000.
- [3] M. E. Crovella and A. Bestavros. Self-similarity in world wide web traffic: Evidence and possible causes. *IEEE/ACM Trans. on Networking*, 5:835–846, 1997.
- [4] John Doyle, Jean Carlson, Steven H. Low, Fernando Paganini, Walter Willinger Glenn Vinnicombe, Jason Hickey, Pablo Parrilo, and Lieven Vandenberghe. Robustness and the internet : Theoretical foundations. Preprint, <http://netlab.caltech.edu/internet/>, 2002.
- [5] R. Johari and D. Tan. End-to-end congestion control for the internet: delays and stability. Technical report, Statistical Laboratory, University of Cambridge, 2000.

- [6] F. P. Kelly. Models for self-managed internet. *Philosophical Transactions of The Royal Society*, A358:2335–2348, 2000.
- [7] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan. Rate control in communication networks: shadow prices, proportional fairness, and stability. *Journal of the Operational Research Society*, 49:237–252, 1998.
- [8] L. Massoulié. Stability of distributed congestion control with heterogeneous feedback delays. Technical Report MSR-TR-2000-111, Microsoft Research, 2000.
- [9] Toru Ohira and Ryusuke Sawatari. Phase transition in a computer network traffic model. *Phys. Rev. E*, 58:193195, 1998.
- [10] Committee on Network Science for Future Army Applications. *Network Science*, page 14. The National Academic Press, 2006.
- [11] Committee on Network Science for Future Army Applications. *Network Science*, page 3. The National Academic Press, 2006.
- [12] Antonis Papachristodoulou, Lun Li, and John C Doyle. Methodological frameworks for largescale network analysis and design. *ACM SIGCOMM Computer Communication Review*, 34:7–20, 2004.
- [13] R.V. Sole and S. Valverde. Information transfer and phase transitions in a model of internet traffic. *Physica A*, 289:595–605, 2001.
- [14] S.H. Strogatz. Exploring complex networks. *Nature*, 410:268–276, 2001.
- [15] Glenn Vinnicombe. On the stability of end-to-end congestion control for the internet. Technical Report CUED/FINFENG/TR.398, Engineering Department, University of Cambridge. <http://www.control.eng.cam.ac.uk/gv/internet/>, 2000.
- [16] W. Willinger, R. Govindan, S. Jamin, V. Paxson, and S. Shenker. Scaling phenomena in the internet: Critically examining criticality. *Proceedings of Natl. Acad. Sci. USA*, 99:2573–2580, 2002.
- [17] W. Willinger, M. S. Taqqu, R. Sherman, and D. V. Wilson. Self-similarity through high-variability: statistical analysis of ethernet lan traffic at the source level. *IEEE/ACM Trans. on Networking*, 5:71–86, 1997.
- [18] Walter Willinger and John Doyle. Robustness and the internet : Design and evolution. Preprint, <http://netlab.caltech.edu/internet/>, 2002.
- [19] Jian Yuan and Kevin Mills. Simulating timescale dynamics of network traffic using homogeneous modeling. *J. Res. Natl. Inst. Stand. Technol.*, 11:227–242, 2006.